

Degree Four Plane Spanners: Simpler and Better

Iyad Kanj¹, Ljubomir Perković¹, and Duru Türkoğlu¹

¹ School of Computing, DePaul University
Chicago, IL, USA
{ikanj,lperkovic,dturkoglu}@cs.depaul.edu

Abstract

Let \mathcal{P} be a set of n points embedded in the plane, and let \mathcal{C} be the complete Euclidean graph whose point set is \mathcal{P} . Each edge in \mathcal{C} between two points p, q is realized as the line segment $[pq]$, and is assigned a weight equal to the Euclidean distance $|pq|$. In this paper, we show how to construct in $\mathcal{O}(n \lg n)$ time a plane spanner of \mathcal{C} of maximum degree at most 4 and stretch factor at most 20. This improves a long sequence of results on the construction of plane spanners of \mathcal{C} . Our result matches the smallest known upper bound of 4 by Bonichon et al. on the maximum degree of plane spanners of \mathcal{C} , while significantly improving their stretch factor upper bound from 156.82 to 20. The construction of our spanner is based on Delaunay triangulations defined with respect to the equilateral-triangle distance, and uses a different approach than that used by Bonichon et al. Our approach leads to a simple and intuitive construction of a well-structured spanner, and reveals useful structural properties of the Delaunay triangulations defined with respect to the equilateral-triangle distance. The structure of the constructed spanner implies that when \mathcal{P} is in convex position, the maximum degree of this spanner is at most 3. Combining the above degree upper bound with the fact that 3 is a lower bound on the maximum degree of any plane spanner of \mathcal{C} when the point-set \mathcal{P} is in convex position, the results in this paper give a tight bound of 3 on the maximum degree of plane spanners of \mathcal{C} for point-sets in convex position.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, I.3.5 Computational Geometry and Object Modeling

Keywords and phrases geometric spanners; plane spanners; bounded degree spanners; Delaunay triangulations

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Let \mathcal{P} be a set of n points embedded in the plane, and let \mathcal{C} be the complete Euclidean graph whose point-set is \mathcal{P} . Each edge in \mathcal{C} between two points p, q is realized as the line segment $[pq]$, and is assigned a weight equal to the Euclidean distance $|pq|$. For any constant $\rho \geq 1$, a subgraph \mathcal{S} of \mathcal{C} is a ρ -spanner of \mathcal{C} if for any two points $p, q \in \mathcal{P}$, there is a path between p and q in \mathcal{S} whose weight is at most $\rho \cdot |pq|$. A subgraph \mathcal{S} of \mathcal{C} is a *spanner* of \mathcal{C} if there exists a constant $\rho \geq 1$ such that \mathcal{S} is a ρ -spanner of \mathcal{C} ; the minimum such ρ is referred to as the *stretch factor* of \mathcal{S} . In this paper, we consider the problem of constructing a plane spanner of \mathcal{C} of small degree and small stretch factor. This problem has received considerable attention, and there is a long list of results on the construction of plane spanners of \mathcal{C} that achieve various trade-offs between the degree and the stretch factor of the spanner.

The problem of constructing a plane spanner of \mathcal{C} was considered as early as the 1980's, if not earlier. Chew [10] proved that the L_1 -Delaunay triangulation of \mathcal{P} , which is the Delaunay triangulation of \mathcal{P} defined with respect to the L_1 -distance, is a $\sqrt{10}$ -spanner of \mathcal{C} . Chew's



© Iyad Kanj, Ljubomir Perković, and Duru Türkoğlu;
licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

result was followed by a series of papers showing that other Delaunay triangulations are plane spanners as well. In 1987, Dobkin *et al.* [13] showed that the classical L_2 -Delaunay triangulation of \mathcal{P} is a plane spanner of stretch factor at most $\frac{\pi(1+\sqrt{5})}{2}$. This bound was subsequently improved by Keil and Gutwin [15] to $\frac{4\pi}{3\sqrt{3}}$. In the meantime, Chew [11] showed that the TD -Delaunay triangulation defined using a distance function based on an equilateral triangle — rather than a square (L_1 -distance) or a circle (L_2 -distance) — is a 2-spanner. This result was generalized by Bose *et al.* [6], who showed that the Delaunay triangulation defined with respect to any convex distance function (*i.e.*, based on a convex shape) is a spanner of \mathcal{C} . The bound on the stretch factor of the L_2 -Delaunay triangulation by Keil and Gutwin stood unchallenged for many years until Xia recently improved the bound to below 2 [20]. Recently as well, Bonichon *et al.* [3] improved Chew’s original bound on the stretch factor of the L_1 -Delaunay triangulation to $\sqrt{4 + 2\sqrt{2}}$, and showed this bound to be tight.

All the Delaunay triangulations mentioned above can have unbounded degree. In recent years, bounded degree plane spanners have been used as the building block of wireless network topologies. Wireless distributed system technologies, such as wireless ad-hoc and sensor networks, are often modeled as proximity graphs in the Euclidean plane. Spanners of proximity graphs represent topologies that can be used for efficient communication. For these applications, in addition to having low stretch factor, spanners are typically required to be plane and have bounded degree, where both requirements are useful for efficient routing [8, 19].

The wireless network applications motivated researchers to turn their attention to minimizing the maximum degree of the plane spanner as well as its stretch factor. It can be readily seen that 3 is a lower bound on the maximum degree of a spanner of \mathcal{C} , because a Hamiltonian path through a set of points arranged in a grid has unbounded stretch factor. Work on bounded degree but not necessarily plane spanners of \mathcal{C} closely followed the above-mentioned work on plane spanners. In a 1992 breakthrough, Salowe [18] proved the existence of spanners of maximum degree 4. The question was then resolved by Das and Heffernan [12] who showed that spanners of maximum degree 3 always exist. The focus in this line of research was to prove the existence of low degree spanners and the techniques developed to do so were not tuned towards constructing spanners that had both low degree and low stretch factor. For example, the bound on the stretch factor of the degree 4 spanner in Salowe [18] is greater than 10^9 , which is far from practical. Furthermore, the bounded-degree spanners shown to exist are not guaranteed to be plane.

Bose *et al.* [7] were the first to show how to extract a subgraph of the L_2 -Delaunay triangulation that is a bounded-degree, plane spanner of \mathcal{C} . The maximum degree and stretch factor they obtained were subsequently improved by Li and Wang [16], by Bose *et al.* [9], and by Kanj and Perković [14] (see Table 1). The approach used in all these results was to extract a bounded degree spanning subgraph of the L_2 -Delaunay triangulation and the main goal was to obtain a bounded-degree plane spanner of \mathcal{C} with the smallest possible stretch factor. In a breakthrough result, Bonichon *et al.* [2] lowered the bound on the maximum degree of a plane spanner from 14 to 6. Instead of using the L_2 -Delaunay triangulation as the starting point of the spanner construction, they used the Delaunay triangulation based on the equilateral-triangle distance, defined originally by Chew [11], and exploited a connection between these Delaunay triangulations and $\frac{1}{2}$ - θ graphs. The plane spanner they constructed also has a small stretch factor of 6. Independently, Bose *et al.* [5] were also able to obtain a plane spanner of maximum degree at most 6, by starting from the L_2 -Delaunay triangulation. Recently, Bonichon *et al.* [4] were able to construct a plane spanner of degree at most 4 and stretch factor at most 156.82. Their construction is based on the Delaunay

Paper	Δ	Stretch factor bound
Bose <i>et al.</i> [7]	27	$(\pi + 1)C_0 \approx 8.27$
Li and Wang [16]	23	$(1 + \pi \sin \frac{\pi}{4})C_0 \approx 6.43$
Bose <i>et al.</i> [9]	17	$(2 + 2\sqrt{3} + \frac{3\pi}{2} + 2\pi \sin(\frac{\pi}{12}))C_0 \approx 23.56$
Kanj and Perković [14]	14	$(1 + \frac{2\pi}{14 \cos(\frac{\pi}{14})})C_0 \approx 2.91$
Bonichon <i>et al.</i> [2]	6	6
Bose <i>et al.</i> [5]	6	$1/(1 - \tan(\pi/7)(1 + 1/\cos(\pi/14)))C_0 \approx 81.66$
Bonichon <i>et al.</i> [4]	4	156.82
This paper	4	20

■ **Table 1** Results on plane spanners with maximum degree bounded by Δ . The constant $C_0 = 1.998$ is the best known upper bound on the stretch factor of the L_2 -Delaunay triangulation [20].

triangulation defined with respect to the L_1 norm. Most of the above spanner constructions can be performed in time $\mathcal{O}(n \lg n)$, where n is the number of points in \mathcal{P} .

In this paper, we present a construction of a plane spanner \mathcal{S} of \mathcal{C} of degree at most 4 and stretch factor at most 20. This result matches the smallest known upper bound of 4 on the maximum degree of the spanner by Bonichon *et al.* [4], while significantly improving their stretch factor bound from 156.82 to 20. Our construction is also simpler and more intuitive. It is based on Delaunay triangulations defined with respect to the equilateral-triangle distance, just like the degree 6 spanner construction used by Bonichon *et al.* [2], which could be viewed as the starting point of our construction. To get down to maximum degree 4, our approach introduces fresh techniques in both the construction and the analysis of the spanner. Unlike the approach in [2], our approach has a bias towards certain edges of the Delaunay triangulation; this bias ensures that the constructed spanner is well structured. To make up for edges not in the spanner, we make use of recursion which, unlike the construction in [4], may have depth not bounded by a constant. To ensure that the recursion is controlled and yields short paths, we aggressively add shortcut edges to the spanner to ensure the existence of paths with specific monotonicity properties, which we refer to as monotone weak paths. Finally, in our analysis we use a new type of distance metric and we also take the extra step of analyzing the stretch factor of our spanner with respect to \mathcal{C} directly, rather than with respect to the underlying Delaunay triangulation.

The structure of the spanner, attained due to biasing certain Delaunay edges, allows us to prove that if the given point-set is in convex position, then the constructed spanner has maximum degree at most 3. Therefore, for any point-set in convex position, there exists a plane spanner of \mathcal{C} of maximum degree at most 3. We also show that 3 is a lower bound on the maximum degree of plane spanners of \mathcal{C} for point-sets in convex position. The preceding implies that 3 is a tight bound on the maximum degree of plane spanners of \mathcal{C} for point-sets in convex position, and this completely and satisfactorily resolves the question about the maximum degree of plane geometric spanners of \mathcal{C} for point-sets in convex position.

2 Preliminaries

Given a set of points \mathcal{P} embedded in the Euclidean plane, we consider the complete weighted graph $\mathcal{C}(\mathcal{P})$, or simply \mathcal{C} , where each edge between any two points $p, q \in \mathcal{P}$ is associated with the line segment $[pq]$, and is assigned a weight equal to the Euclidean distance $|pq|$. Given a subgraph G of \mathcal{C} , we define G to be a *plane* subgraph if the edges of G do not cross

each other, *i.e.*, the line segments associated with the edges of G intersect only at their endpoints. The *maximum degree* of G is the maximum degree (in G) over all points in \mathcal{P} ; we say that a family of graphs has *bounded degree* if there is an integer constant $c \geq 0$ such that every graph in the family has a maximum degree at most c . If graph G is connected, we define the *distance* between any two points $p, q \in \mathcal{P}$, denoted $d_G(p, q)$, to be the weight of a minimum-weight path between p and q in G , where the weight of a path is the sum of the weights of its edges.

Given a constant $\rho \geq 1$, we say that G is a ρ -*spanner* of \mathcal{C} if for any two points $p, q \in \mathcal{P}$, $d_G(p, q) \leq \rho \cdot |pq|$. We also say that a family of geometric graphs, one for every finite set \mathcal{P} of points in the plane, is a *spanner* if there is a constant $\rho \geq 1$ such that every graph $G(\mathcal{P})$ in the family is a ρ -spanner of $\mathcal{C}(\mathcal{P})$; we refer to the minimum such constant ρ as the *stretch factor* of the family. In this paper, the family we consider consists of the spanners obtained by applying our algorithm on all finite point-sets in the plane.

In this paper, we rely on a metric that is different from the Euclidean metric. In order to define this metric, we fix an equilateral triangle with two of its points lying on the x -axis at coordinates $(0, 0)$ and $(1, 0)$, and the third point lying below the x -axis; we use the symbol ∇ to refer to this equilateral triangle. We define a triangle to be a ∇ -*homothet* if it can be obtained through a translation of ∇ followed by a scaling. We define the *triangular metric*, d_∇ , as follows:

► **Definition 1.** For any two points $p, q \in \mathcal{P}$, we define $d_\nabla(p, q)$ to be the side-length of the smallest ∇ -homothet that contains p and q on its boundary; we denote this triangle $\nabla(p, q)$.

It is easy to verify that d_∇ is indeed a metric. In particular, for any two points p, q , we have $d_\nabla(p, q) = 0 \Leftrightarrow p = q$, we have symmetry as in $d_\nabla(p, q) = d_\nabla(q, p)$, and for any third point r , we have the triangle inequality $d_\nabla(p, q) \leq d_\nabla(p, r) + d_\nabla(r, q)$. It is also easy to see that p or q must be a vertex of the triangle $\nabla(p, q)$ and that $|pq| \leq d_\nabla(p, q)$.

Using the triangular metric d_∇ , we define a subgraph \mathcal{D} of \mathcal{C} as follows. For every point $w \in \mathcal{P}$, we partition the space around w into six equiangular cones whose common apex is w , three above and three below the horizontal line passing through w , as illustrated in Figure 1-(a). We denote the middle cone above the horizontal line and the two outer cones below the horizontal line as the *positive* cones of w , and the remaining three cones as the *negative* cones of w . Each point w chooses an edge in each of its three positive cones by selecting the point $v \neq w$ in the cone such that $d_\nabla(w, v)$ is minimum. Assuming that \mathcal{P} is in general position¹, for any two distinct points v, v' in a positive cone of w , we obtain $d_\nabla(w, v) \neq d_\nabla(w, v')$. Let \mathcal{D} be the graph whose vertex-set is \mathcal{P} and whose edge-set is the set of edges selected as described.

We make the following observation regarding the graph \mathcal{D} . The $\frac{1}{2}$ - θ graph of \mathcal{P} is the graph whose point-set is \mathcal{P} , and whose edges are obtained as follows: at each point w , and for each of the three positive cones of apex w , select the edge wv in the cone where v is the point whose projection distance to the angular bisector of the cone is minimum. Bonichon et al. [1] showed that the $\frac{1}{2}$ - θ graph of \mathcal{P} is the same as the *TD-Delaunay triangulation* of \mathcal{P} [11] defined based on the empty triangle property: there is an edge between two point $v, w \in \mathcal{P}$ if there exists a homothet of ∇ containing v and w on its boundary whose interior

¹ \mathcal{P} is in general position if no pair of points $v, w \in \mathcal{P}$ lie on a line parallel to any of the boundary lines defining the six cones. We note that it is always possible to rotate the equilateral triangle that defines the metric d_∇ to ensure that the finite set \mathcal{P} is in general position and so the results in this paper hold for all sets of points and not just for points in general position.

is empty of points of \mathcal{P} . It is easy to see that the $\frac{1}{2}\theta$ graph of \mathcal{P} coincides with the graph \mathcal{D} defined above, and hence with the TD-Delaunay triangulation².

For convenience, we label the positive cones at each point of \mathcal{P} , in clockwise order and starting with the positive cone above the horizontal line, *red*, *green*, and *blue*; we also label the negative cones, in clockwise order and starting with the negative cone below the horizontal line, red, green, and blue. We assign an orientation and a color to the edges of \mathcal{D} by orienting each edge outwards from the point w that selects it and by coloring it red, blue, or green depending on whether the edge lies in the positive red, blue, or green cone of point w , as illustrated in Figure 1-(a). We emphasize that the edge orientations are only used for the purpose of constructing the spanner and proving its desired properties; the final spanner in our construction is an undirected graph obtained by removing edge orientations. In fact, we will abuse terminology and, throughout the paper, use the term *path* to refer to weak paths in \mathcal{D} ; we will always use the term *directed path* when edge orientations are relevant.

We observe that for any point $w \in \mathcal{P}$ there is at most one edge outgoing from w in a positive cone of w , but there can be an unbounded number of edges incoming to w in a negative cone of w , and that in such cases all these edges have the same color as the cone itself (e.g., see Figure 1-(b)). We follow the same approach as Bonichon et al. [2], and identify in each negative cone of point w an edge that plays a key role in the spanner construction:

► **Definition 2.** For any point $w \in \mathcal{P}$, and for each negative cone of w that contains at least one edge incoming to w , let (directed) edge $(v, w) \in \mathcal{D}$ be the edge in the cone such that $d_{\nabla}(v, w)$ is minimum. We define (v, w) to be the *anchor* of w in the cone.


We say that anchors incident to the same point w are *adjacent* if their cones are adjacent. Note that for any two adjacent anchors incident to w , one of the two adjacent anchors must lie in a positive cone of w and must be an anchor of a point other than w .

Consider a negative cone of a point $w \in \mathcal{P}$ containing at least one incoming edge to w in \mathcal{D} . Let $(v_1, w), \dots, (v_k, w) \in \mathcal{D}$, where $k \geq 1$, be all the incoming edges to w that lie in the cone, listed in counterclockwise order, as illustrated in Figure 1-(b), and let (v_j, w) , for some j such that $1 \leq j \leq k$, be the anchor of w in the cone. We call $(v_1, w), \dots, (v_k, w)$ the *fan* of the anchor (v_j, w) . We identify (v_j, w) as the anchor of each edge in the fan. Note that every edge in \mathcal{D} has an anchor which could be itself. We call the first edge (v_1, w) and the last edge (v_k, w) of the fan the *boundary edges* of the anchor (v_j, w) . Note that either one (possibly both) of the boundary edges of an anchor could be the anchor itself. If $k \geq 2$, since \mathcal{D} is a triangulation, it follows that (v_i, v_{i+1}, w) is a triangle in \mathcal{D} , for $i = 1, \dots, k-1$. Hence, v_1, \dots, v_k is a (weak) path in \mathcal{D} between the endpoints v_1 and v_k . We call this path the *canonical path* of w in the designated cone; we also call each edge on this path a *canonical edge* of w . Finally, we refer to the (weak) subpath v_r, \dots, v_s of the canonical path v_1, \dots, v_k of w as the canonical path between v_r and v_s of w . The two *sides* of an edge are the two half-planes defined by the line obtained by extending the edge. We say that a canonical edge e is canonical *on a side* of e if e is a canonical edge of a point that lies on that side of e . Note that a canonical edge can be canonical on both sides.

We state the following easy to verify facts without proof:

► **Lemma 3.** Let (s, t) be a canonical edge of a point w , and let (s', t) be the anchor of (s, t) .
a. The edges (s, w) and (t, w) are in \mathcal{D} .

² A TD-Delaunay triangulation of \mathcal{P} is not necessarily a triangulation of \mathcal{P} as defined traditionally (a triangulation of the convex hull of the set of points). Just as Chew [11] did, we abuse the term *triangulation* because TD-Delaunay triangulations are closely related to classical L_2 -Delaunay triangulations.



■ **Figure 1** (a) To construct graph \mathcal{D} , every point w chooses the shortest edge, according to the d_{∇} distance, in every positive cone. (b) Edge (v_2, w) is the anchor of w in the negative cone shown because it is the shortest edge, according to the d_{∇} distance, among edges incoming to w in the cone; the path v_1, \dots, v_k is the canonical path of anchor (v_2, w) .

- b. The edge (s, w) cannot be a canonical edge on the side containing t .
- c. The edge (t, w) is not an anchor.
- d. The edge (s, t) is a boundary edge of its anchor (s', t) .


3 Monotone (weak) paths

We define next a type of path in \mathcal{D} that generalizes canonical paths and that will be a key tool in our construction. We give two equivalent definitions of such a path; proving the equivalence between these two definitions is easy.

► **Definition 4.** Let v be a point lying in the positive cone of u whose color is c . A (weak) path in \mathcal{D} between u and v is *monotone* if the path is bi-colored, with c being one of the colors, and the path satisfies the two equivalent properties:

- After reversing the direction of all edges not colored c , the path is directed from u to v .
- No two consecutive edges of the path lie in neighboring cones of the shared endpoint.

The key property of a monotone path between u and v is that its length can be bounded by twice the side-length of $\nabla(u, v)$, i.e., by $2d_{\nabla}(u, v)$. This follows from a stronger insight



■ **Figure 2** (a) If point v lies in the positive cone of point u , and the vertices of $\nabla(u, v)$, u, y, z , are colored green, red, and blue, respectively, then $d_{\nabla}(u, v) = |zu|$ and $\delta_{\nabla}^{blue}(u, v) = |zv|$. (b) and (c) P is a monotone path between u and v with edges colored green or red. The projection onto zu (resp. zy) of (u, p_2) , (p_1, p_2) , and (p_1, v) do not overlap and are contained within $[zu]$ (resp. $[zv]$).

which we develop next. To facilitate our discussion, we label the vertices of a ∇ -homothet *green*, *blue*, and *red*, in clockwise order starting from the upper left vertex. Let v be a point lying in a positive cone of u of color c_1 . With u being the vertex of $\nabla(u, v)$ of color c_1 , let the remaining two vertices of $\nabla(u, v)$ be y and z , and let c_2 be the color of y and c_3 be the color of z .

► **Definition 5.** We define (refer to Figure 2 where $c_1 = \text{green}$, $c_2 = \text{red}$, and $c_3 = \text{blue}$):

1. $\delta_{\nabla}^{c_2}(u, v) = \delta_{\nabla}^{c_2}(v, u) = |yv|$.
2. $\delta_{\nabla}^{c_3}(u, v) = \delta_{\nabla}^{c_3}(v, u) = |zv|$.
3. $\delta_{\nabla}^{\min}(u, v) = \delta_{\nabla}^{\min}(v, u) = \min\{\delta_{\nabla}^{c_2}(u, v), \delta_{\nabla}^{c_3}(u, v)\}$.

Let P be a monotone path in \mathcal{D} between u and v , and let the edges of P be colored c_1 or c_2 . We define, using the lines zv and zu as axes of a coordinate system of the Euclidean plane, the *projection onto zv and zu* (see Figure 2-(b)). In the following lemma, we use this projection to map the edges of P to upper bound the length of P :

► **Lemma 6.** Let v be a point lying in the positive cone of u of color c_1 . For any monotone path P between u and v whose edges are colored c_1 or c_2 (refer to Figure 2 where $c_1 = \text{green}$, $c_2 = \text{red}$, and $c_3 = \text{blue}$):

- a. The projections of all edges of P onto zu (resp., zv) do not overlap and are contained within the segment $[zu]$ (resp., $[zv]$); see Figures 2-(b) and 2-(c).
- b. If (p, q) is an edge of P colored c_1 (resp., c_2) then the projection onto zu (resp., zv) of (p, q) has length $d_{\nabla}(p, q) \geq |pq|$.
- c. The sum of the lengths of the edges of P colored c_1 is at most $d_{\nabla}(u, v) = |zu|$.
- d. The sum of the lengths of the edges of P colored c_2 is at most $\delta_{\nabla}^{c_3}(u, v) = |zv|$.
- e. The length of P is at most $2d_{\nabla}(u, v)$.

Proof. For part (a), we consider the coordinates of the points of P in the coordinate system of the Euclidean plane defined by using the lines zv and zu as axes. When visiting the points of P in the order in which they appear on P , the coordinates of the points along the zu (resp., zv) axis form a monotonic sequence (decreasing or increasing) between the coordinates of u and z (resp. z and v), and part (a) follows. Since zu is parallel to an edge of $\nabla(p, q)$, and hence the projection of (p, q) onto zu has length $d_{\nabla}(p, q)$, part (b) follows. Parts (c) and (d) follow from parts (a) and (b), and part (e) follows from parts (c) and (d). ◀

The following lemma, illustrated in Figure 3-(a), is an insight implicit in Lemma 2 of [2] that is implied by Lemma 6.

► **Lemma 7.** For any two edges (v, w) and (u, w) that lie in the same fan:

- a. The canonical path P between v and u is monotone.
- b. The sum of the lengths of all monochromatic edges on P is at most $d_{\nabla}(v, u)$.
- c. The length of the canonical path P between v to u is at most $2d_{\nabla}(v, u)$.

Proof. For part (a), we assume, without loss of generality, that w lies in the red positive cones of v and of u . We then observe that for every point p on P , (p, w) is an edge in \mathcal{D} . Therefore, every edge of P must lie in the blue or green positive cones of its tail, and thus path P is bi-colored. Furthermore, since \mathcal{D} is planar, at every intermediate point p of P , the two edges of P incident to p must lie in non-adjacent cones. The canonical path P between v and u is thus monotone. Hence, parts (b) and (c) follow by Lemma 6. ◀




Figure 3 (a) Illustration of Lemma 7. (b) w is incident to an anchor in every cone. In step 1, both of the blue anchors are added. In step 2, on the other hand, no more than one white anchor is added below the horizontal line through w and no more than one white anchor is added above.

4 The Spanner

In this section, we describe the construction of a plane spanner of \mathcal{C} of maximum degree at most 4 and stretch factor at most 20. In our construction, we will bias blue – positive and negative – cones and edges. This bias results in a spanner satisfying structural properties that allow us to prove the desired upper bounds on the spanner degree and stretch factor. This bias also ensures that the spanner has maximum degree at most 3 when the point-set \mathcal{P} is in convex position. We find it convenient to refer to the four non-blue cones, as well as all the red and green edges, as *white*, as illustrated in Figure 3-(b), and introduce some terminologies.

If e is a canonical edge of a point w that lies in a white (resp., blue) cone of w , we say that e is a canonical edge *in a white* (resp., *blue*) *cone*. We note that a canonical edge could be in a white cone of one point and in a blue cone of another. Given a white anchor (v, w) , the ray starting from w extending (v, w) partitions the (white) negative cone of w containing v into two *sides*: we refer to the side of the cone that is adjacent to a blue cone as the *blue side*, and we refer to the other side that is adjacent to a white cone as the *white side*. We say that an edge (u, w) in the fan of (v, w) is on the *white side* (resp. *blue side*) if it is on the white side (resp. blue side) of (v, w) .

The following describes the construction of the spanner \mathcal{S} of \mathcal{C} . The construction is based on the underlying triangulation \mathcal{D} of \mathcal{C} . We start by constructing a degree-4 anchor subgraph \mathcal{A} of \mathcal{S} that includes all blue anchors. We then finalize \mathcal{S} by adding some white canonical edges and shortcuts to ensure the reconstructibility of all canonical paths.


Constructing the anchor subgraph \mathcal{A} of \mathcal{S} :

1. We add to \mathcal{A} every blue anchor (of \mathcal{D}).
2. In increasing order of length with respect to the metric d_{∇} , for every white anchor a , we add a to \mathcal{A} if no white anchor adjacent to a is already in \mathcal{A} .

Reconstructing canonical paths in blue cones (see Figure 4-(a)):

3. We add to \mathcal{S} every (white) canonical edge in a blue cone if the edge is not in \mathcal{A} .
4. For every pair of canonical edges $(p, q), (r, q)$ in a blue cone such that $(p, q), (r, q) \in \mathcal{S} \setminus \mathcal{A}$, we add to \mathcal{S} the shortcut edge (p, r) , color it white, and remove (p, q) and (r, q) from \mathcal{S} .

Reconstructing canonical paths in white cones (see Figure 4-(b)):



■ **Figure 4** (a) In step 3, white canonical edges of w in the negative blue cone of w are added to \mathcal{S} if not in \mathcal{A} already; in step 4, any pair of canonical edges of w added in step 3 that are incoming at the same point are replaced by a shortcut between the outgoing endpoints ((p_6, p_5) and (p_4, p_5) replaced by shortcut (p_6, p_4) and (p_3, p_2) and (p_1, p_2) replaced by shortcut (p_3, p_1) .) (b) Shortcut edges (p_3, p_0) and (p_7, p_4) are added to \mathcal{S} in step 6; edges (p_7, p_6) and (p_3, p_2) are not in \mathcal{S} unless they are anchors in \mathcal{A} .

5. We add to \mathcal{S} every white canonical edge that is on the white side of its (white) anchor, but only if its anchor is not in \mathcal{A} .
6. For every white anchor (v, w) and its boundary edge $(u, w) \neq (v, w)$ on the white side, let P be the canonical path $(u = p_0, p_1, \dots, p_k = v)$. We apply the following procedure at a current point p_i starting with $i = 0$ and stopping when $i = k$:
 - a. If the canonical edge (p_{i+1}, p_i) is white, we skip this edge and set i to $i + 1$;
 - b. Otherwise, (p_i, p_{i+1}) must be blue. Let $j > i$ be the largest index of a point on P such that the line segment $[p_i p_j]$ does not intersect the canonical path from p_i to p_j (except at p_i and p_j). We add the shortcut (p_j, p_i) to \mathcal{S} and color it white; we remove the (white) canonical edge (p_j, p_{j-1}) from \mathcal{S} if $(p_j, p_{j-1}) \in \mathcal{S} \setminus \mathcal{A}$; and we set i to j .

In the following section we prove that this algorithm yields a plane spanner of maximum degree at most 4 and stretch factor at most 20. We provide here a high-level overview of our arguments.

To show planarity, we note that the underlying graph \mathcal{D} is planar and that the only edges of \mathcal{S} not in \mathcal{D} are the shortcut edges added in steps 4 and 6.b. We will prove in Lemmas 10 and 11 that each such edge does not intersect any other edge of \mathcal{S} . For the degree upper bound, we note that the first two steps of the algorithm yield the subgraph \mathcal{A} of maximum degree at most 4. In the remaining steps, we carefully add additional edges, whether canonical edges or shortcuts of canonical paths. To prove the degree bound, we develop a charging argument that assigns each edge of \mathcal{S} to a cone at each endpoint and show, in Lemma 13, that no more than 4 cones are charged at every point.

Note that step 1 adds all blue anchors to \mathcal{S} . For white anchors, we cannot include all of them in \mathcal{S} because of the degree bound. However, by considering anchors in an increasing order of their length in step 2, we can ensure that a white anchor is not added to \mathcal{S} only if an adjacent shorter white anchor was added to \mathcal{S} . This ordering is crucial because it allows us to use short anchors in \mathcal{S} to reconstruct longer anchors not in \mathcal{S} .

In steps 3 and 4, we add to \mathcal{S} all (white) canonical edges in blue cones, except for some consecutive pairs of canonical edges that are replaced with shortcut edges. Therefore, steps 1 through 4 of the algorithm ensure that every canonical path in a blue cone has been

reconstructed. In particular, since all blue anchors are in \mathcal{S} , every blue boundary edge, and therefore every blue canonical edge, has been reconstructed.

Consider now a canonical path P in a white negative cone of w , and the two subpaths P_w and P_b of P on the white side and the blue side of this cone, respectively. In step 6, we take shortcuts to obtain, together with step 5, a white monotone path to ensure that there is a short path reconstructing P_w . As for P_b , we argued above that the blue edges of P_b are reconstructible. As for the white edges of P_b , we note that they are on the white side of their anchors, and hence, the argument used to reconstruct P_w above applies to them as well. The above, once again with step 5, ensures the reconstructibility of P_b , and hence of P .

5 Properties of the Spanner

In this section, we prove the three properties of the spanner obtained using our algorithm: planarity, the maximum degree upper bound of 4, and the stretch factor bound of 20. We start with the following justification for coloring shortcut edges added in step 6 white:

► **Lemma 8.** *For every shortcut edge (p_j, p_i) added to \mathcal{S} in step 6, p_j and $v = p_k$ both lie in the same negative white cone of p_i , and they both lie in the same negative white cone of p_{j-1} .*

Proof. Because $d_{\nabla}(w, p_k) < d_{\nabla}(w, p_i)$ and $(p_i, w) \in \mathcal{D}$, p_k must lie in a negative white cone of p_i . The lemma thus holds if $j = k$. Otherwise, by the choice of p_j , p_j must lie on the same side of line $p_i p_k$ as point w ; again, because $(p_i, w) \in \mathcal{D}$, p_j must lie in a (negative) white cone of p_i that also contains p_k . Similar arguments apply to p_{j-1} . ◀

Next, we show that \mathcal{S} is plane. We first need the following definition and lemma.

► **Definition 9.** An edge $(u, w) \in \mathcal{D}$ is *uncrossed* if no shortcut in \mathcal{S} crosses (u, w) .

► **Lemma 10.** *All anchors, all canonical edges, and all boundary edges are uncrossed.*

Proof. Let (p, r) be a shortcut that was added in step 4 of the spanner construction, let (p, q) and (r, q) be the pair of canonical edges in the blue cone as described in step 4, and let w be the apex of this blue cone. It is easy to verify that $(q, w) \in \mathcal{D}$ is the only edge in \mathcal{D} that (p, r) crosses, and that (q, w) is not a boundary edge, a canonical edge, or an anchor. Next, consider a shortcut (p_j, p_i) that was added in step 6 of the spanner construction, and let (v, w) be the white anchor and $p_{i+1}, p_{i+2}, \dots, p_{j-1}$ be the points on the canonical path between p_i and p_j as described in step 6. Again, it is easy to verify that $(p_{i+1}, w), (p_{i+2}, w), \dots, (p_{j-1}, w)$ are the only edges in \mathcal{D} that the shortcut (p_j, p_i) crosses, and that none of them is a boundary edge, a canonical edge, or an anchor. ◀

► **Lemma 11.** *The subgraph \mathcal{S} is a plane subgraph of \mathcal{C} .*

Proof. Let $\mathcal{S}_1 = \mathcal{D} \cap \mathcal{S}$ and $\mathcal{S}_2 = \mathcal{S} \setminus \mathcal{S}_1$. Note that \mathcal{S}_1 consists of \mathcal{A} plus the canonical edges added in steps (3) or (5) that are kept after steps (4) and (6), and that \mathcal{S}_2 consists only of the shortcuts added in steps (4) and (6). Since \mathcal{S}_1 is a subgraph of \mathcal{D} , \mathcal{S}_1 is plane. By Lemma 10, all the edges in \mathcal{S}_1 are uncrossed, i.e., no shortcut (edge in \mathcal{S}_2) crosses an edge in \mathcal{S}_1 . To conclude the proof, we show that no two edges in \mathcal{S}_2 cross either. Observe that any two shortcuts connect pairs of endpoints of canonical paths that either belong to different fans, and in which case they cannot cross, or belong to the same fan. In the latter case, the shortcuts do not cross because shortcuts always connect the endpoints of non-overlapping canonical paths. ◀

To facilitate the discussion in the proof of the degree upper bound, we refer to the two adjacent white cones above (resp., below) the horizontal line through a point $p \in \mathcal{P}$ as the *upper* (resp., *lower*) *white sector* of p ; we also refer to the two blue cones at p as the *left* and *right blue sectors* of p . We develop next a charging scheme and use it to show that, for each point p , each edge incident to p in \mathcal{S} can be mapped in a one-to-one fashion to one of the four sectors at p . To describe the charging scheme for every edge $e \in \mathcal{S}$ and for every endpoint p of e , we define $\sigma(e, p)$ to be the sector of p that contains e . Also for a point p , we denote by LB_p , RB_p , UW_p , and LW_p , the left blue, the right blue, the upper white, and the lower white sectors of p respectively. We describe in the table below the charging scheme for every edge $e = (x, y) \in \mathcal{S}$ based on which step of the construction e is added to \mathcal{S} .

Step	Classification of $e = (x, y)$	Charge at x	Charge at y
1	Blue anchor in \mathcal{A}	$\sigma(e, x) = LB_x$	$\sigma(e, y) = RB_y$
2	White anchor in \mathcal{A}	$\sigma(e, x) = UW_x$ or LW_x	$\sigma(e, y) = LW_y$ or UW_y
3	White canonical edge in a blue cone	$\sigma(e, x) = UW_x$ or LW_x	LB_y
4	Shortcut in a blue cone	$\sigma(e, x) = UW_x$ or LW_x	$\sigma(e, y) = LW_y$ or UW_y
5	White canonical edge in a white cone	RB_x	$\sigma(e, y) = UW_y$ or LW_y
6	Shortcut in a white cone	RB_x	$\sigma(e, y) = UW_y$ or LW_y

► **Lemma 12.** *Let $p \in \mathcal{P}$. The charging scheme above charges each edge incident to p in \mathcal{S} to one of the sectors at p such that each sector is charged with at most one edge.*

Proof. Let $p \in \mathcal{P}$. We will show that the charging scheme above charges every edge incident to p in \mathcal{S} to one of the four sectors at p in a one-to-one fashion.

First, consider the left blue sector of p . Observe that LB_p could potentially be charged in each of the following two situations:

- i. In step (1) by a blue anchor $(x = p, y) \in \mathcal{A}$; and
- ii. in step (3) by a white canonical edge $(x, y = p)$ in a blue cone.

First, observe that (i) and (ii) cannot apply simultaneously by part (c) of Lemma 3. Second, observe that LB_p cannot be charged twice according to situation (ii) because this means that there are two white canonical edges incoming to p in a blue cone, which would imply that step (4) of the spanner construction applies, and both incoming edges to p would be removed from \mathcal{S} . It follows that LB_p is charged at most once.

Next, we consider the right blue sector of p . Observe that RB_p could potentially be charged in each of the following three situations:

- i. In step (1) by a blue anchor $(x, y = p) \in \mathcal{A}$;
- ii. in step (5) by a white canonical edge $(x = p, y)$ in a white cone; and
- iii. in step (6) by a shortcut $(x = p, y)$ in a white cone.

First, observe that in both situations (ii) and (iii), we know that there is a white canonical edge (p, q) in a white cone with apex w , that is (p, w) and (q, w) are white edges in \mathcal{D} ; in (ii) (p, q) is (x, y) ; and in (iii) (p, q) is the (last) edge of the canonical path between $x = p$ and y , in step (6) of the spanner construction, by the choice of point x in this step.

The existence of this white canonical edge implies that RB_p must be empty of edges of \mathcal{D} , that (i) cannot apply when either (ii) or (iii) applies.

Second, by part (b) of Lemma 3, (p, w) is not a canonical edge on the side that contains q , i.e., (p, w) cannot be a canonical edge in a white cone.

Therefore, (p, q) is the only white canonical edge in a white cone that is outgoing from p , and hence, RB_p cannot be charged twice according to (ii), or twice according to (iii).

Finally, (ii) and (iii) cannot apply simultaneously because step (6) of the construction removes this white canonical edge (p, q) when it adds the shortcut (x, y) .

It follows that RB_p is charged at most once.

Finally, we consider the upper sector of p , and note that the arguments for the lower sector of p follow similarly. Observe that UW_p could potentially be charged in each of the following five situations:

- i. In step (2) by a white anchor $(x, y) \in \mathcal{A}$, $x = p$ or $y = p$;
- ii. in step (3) by a white canonical edge $(x = p, y)$ in a blue cone;
- iii. in step (4) by a shortcut (x, y) in a blue cone, $x = p$ or $y = p$;
- iv. in step (5) by a white canonical edge $(x, y = p)$ in a white cone; and
- v. in step (6) by a shortcut $(x, y = p)$ in a white cone.

First, since step (2) of the construction disallows adjacent anchors to be in \mathcal{A} , UW_p cannot be charged twice according to (i).

Second, we show that in both situations (ii) and (iii), none of the other situations (i), (iv), or (v) apply. To show this, we observe that there exists a white canonical edge $(p, q) \notin \mathcal{A}$ in a blue cone with apex w , that is (p, w) and (q, w) are blue edges in \mathcal{D} ; in (ii) (p, q) is (x, y) and in (iii) (p, q) is the canonical edge incident to p that step (4) of the construction removes after adding the shortcut (x, y) .

The existence of this white canonical edge (p, q) implies that the negative white cone of UW_p must be empty of edges of \mathcal{D} , hence, (iv) cannot apply.

By part (b) of Lemma 3, since (p, w) is not a canonical edge on the side that contains q , (v) cannot apply.

Also, since $(p, q) \notin \mathcal{A}$ is the only outgoing edge in UW_p , (i) cannot apply either.

Next, observe that UW_p cannot be charged twice according to (ii) or according to (iii) simply because (p, q) is the only canonical edge in a blue cone that lies in UW_p .

Finally, observe that (ii) and (iii) cannot apply simultaneously because as described above step (4) of the construction removes this canonical edge (p, q) .

Assuming that (ii) and (iii) do not apply, we analyze situation (iv). In this case, we know that the white canonical edge (x, p) is in a white cone with apex w , that is (x, w) and (p, w) are white edges in \mathcal{D} .

Since step (5) of the construction ensures that the anchor of the canonical edge (x, p) is not in \mathcal{A} , and also since by part (c) of Lemma 3, (p, w) is not anchor, thus not in \mathcal{A} , we conclude that (i) does not apply.

Next, since (x, p) is a white and not a blue canonical edge for the fan that contains (x, w) and (p, w) , step (6) of the construction ensures that (v) does not apply.

Finally, observe that UW_p cannot be charged twice according to (iv) simply because (x, p) is the only canonical edge in a white cone that lies in UW_p .

For the remaining cases, we analyze situation (v). In this case, there is a blue canonical edge (p, q) in a white cone with apex w , that is (p, w) and (q, w) are white edges in \mathcal{D} . The existence of this blue canonical edge implies that the white cone contained in UW_p is empty of edges of \mathcal{D} , and because step (6) adds a shortcut incoming at p , we know that the only outgoing white edge $(p, w) \in UW_p$ is not an anchor, thus (i) cannot apply. Finally, we cover all the cases by observing that UW_p cannot be charged twice according to (v) simply because step (6) of the construction adds at most one shortcut for any blue canonical edge, and because (p, q) is the only blue (canonical) edge outgoing from p .

Therefore, each of UW_p and LW_p is charged at most once.

It follows from the above that each of the four sectors at p is charged with at most one edge incident to p in \mathcal{S} . This completes the proof. \blacktriangleleft

► **Lemma 13.** *The maximum degree of \mathcal{S} is at most 4.*

Proof. The statement of the lemma directly follows from Lemma 12 because there are four sectors at each point $p \in \mathcal{P}$. \blacktriangleleft

The remainder of this section is devoted to showing that the stretch factor of \mathcal{S} is at most 20. We do so by first proving a sequence of lemmas that derive upper bounds on the distance in \mathcal{S} between the endpoints of different types of edges in \mathcal{D} ; we then use these lemmas to derive a bound on the stretch factor of \mathcal{S} .

► **Lemma 14.** *For any uncrossed blue edge $(u, w) \in \mathcal{D}$, $d_{\mathcal{S}}(u, w) \leq 3d_{\nabla}(u, w)$.*

Proof. Let (v, w) be the blue anchor of the blue edge (u, w) . In step (1) of the algorithm, we add all the blue anchors in \mathcal{S} , and thus $(v, w) \in \mathcal{S}$. Also, in step (3) of the construction, we add in \mathcal{S} all the canonical edges in blue cones except that, in step (4), we substitute some pairs of these canonical edges with shortcuts. Since (u, w) is uncrossed, these canonical edges and shortcuts provide a path for connecting v and u . Using the triangle inequality, this path that includes the shortcuts is not longer than the canonical path between v and u . Hence, in the worst case, we may assume, without loss of generality, that the path connecting v and u consists only of canonical edges on the canonical path. This canonical path plus the anchor constitutes a path between u and w . By Lemma 7, the length of this canonical path is bounded by $2d_{\nabla}(v, u) \leq 2d_{\nabla}(u, w)$. We also have that $|vw| \leq d_{\nabla}(u, w)$. Consequently, the length of this path is bounded by $d_{\nabla}(u, w)$ (anchor) plus $2d_{\nabla}(u, w)$ (canonical path). It follows that $d_{\mathcal{S}}(u, w) \leq 3d_{\nabla}(u, w)$. \blacktriangleleft

► **Lemma 15.** *For any white anchor (v, w) and any uncrossed white edge $(u, w) \in \mathcal{D}$ that lies on the white side of (v, w) , $d_{\mathcal{S}}(v, u) \leq d_{\nabla}(v, u) + \delta_{\nabla}^{\text{blue}}(v, u) \leq 2d_{\nabla}(v, u)$. Furthermore, if $(v, w) \in \mathcal{S}$, then $d_{\mathcal{S}}(u, w) \leq d_{\nabla}(u, w) + \delta_{\nabla}^{\text{blue}}(u, w) \leq 2d_{\nabla}(u, w)$.*

Proof. We describe below how to construct a white monotone path in \mathcal{S} between u and v . If $(v, w) \in \mathcal{S}$, we extend this path to a white monotone path between u and w . Then, we obtain the desired bounds using Lemma 6.

To describe the white monotone path between u and v , we consider the uncrossed edges of \mathcal{D} on the fan of (v, w) whose endpoints lie on the canonical path between v and u . We observe that shortcuts and white canonical edges connect the (distinct) endpoints of those uncrossed edges, and that they form a white monotone path between u and v because at each point the white edges of the path incident to the point lie on opposite sides of the horizontal line through the point. We call this white monotone path the *white monotone connection* between v and u . We know that all of the shortcuts on this white monotone connection are in \mathcal{S} and even though some of the white canonical edges may not be in \mathcal{S} , we know, for all such white canonical edges, that we have their anchors in \mathcal{S} . For each such white canonical edge (s, t) , we recursively expand the current white monotone path by including the anchor (r, t) of (s, t) and by including the white monotone connection between r and s . We observe that the path obtained after the expansion of (s, t) continues to be a white monotone path. Therefore, by recursively expanding this path for white canonical edges that are not in \mathcal{S} , we obtain a white monotone path between v and u . Furthermore, if $(v, w) \in \mathcal{S}$, we expand this path to include the white anchor (v, w) while preserving its monotonicity. \blacktriangleleft

► **Lemma 16.** *For any white anchor (v, w) and any white edge $(u, w) \in \mathcal{D}$ that lies on the blue side of (v, w) , $d_{\mathcal{S}}(v, u) \leq 5d_{\nabla}(v, u)$. Furthermore, if $(v, w) \in \mathcal{S}$, then $d_{\mathcal{S}}(u, w) \leq 6d_{\nabla}(u, w)$.*

Proof. The canonical path from v to u consists of blue and white canonical edges. The total length of the blue canonical edges does not exceed $d_{\nabla}(v, u)$, and the total length of the white canonical edges does not exceed $d_{\nabla}(v, u)$ by Lemma 7. By Lemma 10, we know that all of these canonical edges are uncrossed. By Lemma 14, the total length of the paths needed to reconstruct these blue canonical edges can be bounded by $3d_{\nabla}(v, u)$. Also, since either the white canonical edges themselves or their anchors are in \mathcal{S} , the total length of the white canonical edges can be bounded by $2d_{\nabla}(v, u)$ by Lemma 15. Therefore, $d_{\mathcal{S}}(v, u)$ can be bounded by $5d_{\nabla}(v, u)$ for the edge (v, u) as stated. Furthermore, if $(v, w) \in \mathcal{S}$, $d_{\mathcal{S}}(u, w)$ can be bounded by $5d_{\nabla}(v, u) + d_{\nabla}(v, w)$, which in turn is bounded by $6d_{\nabla}(u, w)$. ◀

► **Definition 17.** For any two points $p, q \in \mathcal{P}$ such that p lies in a white cone of q , we define $\delta_{\nabla}^{\text{white}}(p, q) = \delta_{\nabla}^{\text{white}}(q, p) = d_{\nabla}(p, q) - \delta_{\nabla}^{\text{blue}}(p, q)$.

► **Lemma 18.** *For any two white edges $(v, w), (u, w) \in \mathcal{D}$ that both lie in the same negative cone of w , if u lies in a positive white cone of v , then $\delta_{\nabla}^{\text{blue}}(u, w) = \delta_{\nabla}^{\text{blue}}(v, w) + d_{\nabla}(v, u)$ and $d_{\nabla}(u, w) = d_{\nabla}(v, w) + \delta_{\nabla}^{\text{blue}}(v, u)$.*


Proof. Since $(v, w) \in \mathcal{D}$, we observe that u must lie in the positive white cone of v that does not contain w . The rest follows directly from the definition of $\delta_{\nabla}^{\text{blue}}$. ◀

► **Lemma 19.** *For any white anchor (v, w) , $d_{\mathcal{S}}(v, w) \leq 9d_{\nabla}(v, w)$. Furthermore, for any uncrossed white edge (u, w) in the fan of (v, w) , we have $d_{\mathcal{S}}(u, w) \leq 9d_{\nabla}(u, w) + \delta_{\nabla}^{\text{blue}}(u, w)$ if (u, w) lies on the white side of (v, w) , and $d_{\mathcal{S}}(u, w) \leq 9d_{\nabla}(u, w)$ otherwise.*

Proof. If $(v, w) \in \mathcal{S}$, then clearly $d_{\mathcal{S}}(v, w) \leq d_{\nabla}(v, w)$. As for any uncrossed edge (u, w) in the fan, by Lemma 15, we get a bound of $2d_{\nabla}(u, w)$ on $d_{\mathcal{S}}(u, w)$ if (u, w) lies on the white side of (v, w) , and by Lemma 16, we get a bound of $6d_{\nabla}(u, w)$ on $d_{\mathcal{S}}(u, w)$ if (u, w) lies on the blue side of (v, w) . Then, we consider $(v, w) \notin \mathcal{S}$ and analyze two cases: (v, w) was not added in \mathcal{A} because of an adjacent anchor at w , or because of an adjacent anchor at v .

If (v, w) was not added in \mathcal{A} because of an adjacent (white) anchor at w , let (w, w') be that anchor (see Figure 5-(a)). By our construction of \mathcal{A} , we know that (w, w') must be shorter than (v, w) , i.e., $d_{\nabla}(w, w') < d_{\nabla}(v, w)$. Therefore, v lies in the positive blue cone of w' , and hence, there must be an outgoing blue edge at w' . Let (w', u') be that blue edge; then, (u', w) must be a white boundary edge of (v, w) , and possibly $u' = v$. Using the fact that u' lies in the positive blue cone of w' , it is easy to verify that $d_{\nabla}(v, u') \leq \delta_{\nabla}^{\text{white}}(v, w) \leq d_{\nabla}(v, w)$. Similarly, using the fact that u' lies in a positive white cone of v , it is easy to verify that $d_{\nabla}(w', u') \leq d_{\nabla}(v, w) + \delta_{\nabla}^{\text{white}}(v, w) \leq 2d_{\nabla}(v, w)$. Following the path from w to w' to u' to v , we bound $d_{\mathcal{S}}(v, w)$ by $d_{\nabla}(w, w')$ for the edge (w, w') , by $3d_{\nabla}(w', u')$ for the edge (w', u') using Lemmas 10 and 14, and by $2d_{\nabla}(v, u')$ for the edge (v, u') using Lemmas 10 and 15. Also, using the above inequalities $d_{\nabla}(w, w') < d_{\nabla}(v, w)$, $d_{\nabla}(v, u') \leq d_{\nabla}(v, w)$, and $d_{\nabla}(w', u') \leq 2d_{\nabla}(v, w)$, we get the desired upper bound $d_{\mathcal{S}}(v, w) \leq 9d_{\nabla}(v, w)$.

Next, we consider the uncrossed white edges. For any uncrossed edge (u, w) on the white side of the anchor, the bound on $d_{\mathcal{S}}(v, w)$ applies directly to $d_{\mathcal{S}}(u, w)$ because the path between v and w already connects u and w . Also, since $d_{\nabla}(v, w) \leq d_{\nabla}(u, w)$, we immediately get $d_{\mathcal{S}}(u, w) \leq 9d_{\nabla}(u, w)$. As for any white edge (u, w) on the blue side of the anchor, we start by observing that $\delta_{\nabla}^{\text{white}}(v, w) \leq d_{\nabla}(u, w) - d_{\nabla}(v, u)$, and that $d_{\nabla}(v, w) \leq d_{\nabla}(u, w)$. Also, using the above inequalities $d_{\nabla}(w, w') < d_{\nabla}(v, w)$, and $d_{\nabla}(v, u') \leq \delta_{\nabla}^{\text{white}}(v, w)$, and $d_{\nabla}(w', u') \leq d_{\nabla}(v, w) + \delta_{\nabla}^{\text{white}}(v, w)$, we obtain the inequalities $d_{\nabla}(w, w') \leq d_{\nabla}(u, w)$, and



■ **Figure 5** Illustrations of the proof of Lemma 19. (a) The case when $(v, w) \notin \mathcal{A}$ because a shorter adjacent anchor (w, w') was added first. Edge (w', u') is a blue boundary edge and there is a white monotone path from u' to v in \mathcal{S} . (b) The case when $(v, w) \notin \mathcal{A}$ because a shorter adjacent anchor (v', v) was added first. Edge (w, u') is a blue boundary edge and (u', v) is a white boundary edge on the white side of its cone and there is a white monotone path between u' and v' in \mathcal{S} .

$d_{\nabla}(v, u') \leq d_{\nabla}(u, w) - d_{\nabla}(v, u)$, and $d_{\nabla}(w', u') \leq 2d_{\nabla}(u, w) - d_{\nabla}(v, u)$. Then, following the above path from w to v and extending it to u using the canonical edges, we get $d_{\mathcal{S}}(u, w) \leq d_{\nabla}(w, w') + 3(d_{\nabla}(w', u') + d_{\nabla}(v, u)) + 2(d_{\nabla}(v, u') + d_{\nabla}(v, u))$, which is then bounded by $d_{\mathcal{S}}(u, w) \leq d_{\nabla}(u, w) + 6d_{\nabla}(u, w) + 2d_{\nabla}(u, w) = 9d_{\nabla}(u, w)$.

In the case when (v, w) was not added in \mathcal{A} because of an adjacent (white) anchor at v , (see Figure 5-(b)) let (v', v) be that anchor. Similarly, we know that (v', v) must be shorter than (v, w) , i.e., $d_{\nabla}(v', v) < d_{\nabla}(v, w)$. Therefore, we know that there must be an outgoing blue edge at w . Let (w, u') be that blue edge; then, (u', v) must be a white boundary edge of (v', v) , possibly $u' = v'$. Once again, it is easy to see that the length of the edge (v', u') cannot exceed the length on the white anchor (v, w) and that the length of the blue edge (w, u') cannot exceed twice the length of the white anchor (v, w) , i.e., $d_{\nabla}(v', u') \leq d_{\nabla}(v, w)$ and $d_{\nabla}(w, u') \leq 2d_{\nabla}(v, w)$. Following the path from w to u' to v' to v , we bound $d_{\mathcal{S}}(v, w)$ by $3d_{\nabla}(w, u') + 2d_{\nabla}(v', u') + d_{\nabla}(v', v)$ using Lemmas 10, 14 and 15. Using the above inequalities, we get the desired upper bound $d_{\mathcal{S}}(v, w) \leq 9d_{\nabla}(v, w)$. As for the uncrossed white edges in this case, we first note that the anchor (v, w) itself is the boundary edge on the blue side and that there are no white edges on the blue side, thus nothing to prove on the blue side. As for any uncrossed white edge (u, w) on the white side of (v, w) , we follow the path from w to v as above and to u using the path described in Lemma 15, which also bounds the last step by $d_{\nabla}(v, u) + \delta_{\nabla}^{\text{blue}}(v, u)$. We get an upper bound $d_{\mathcal{S}}(u, w) \leq 9d_{\nabla}(v, w) + d_{\nabla}(v, u) + \delta_{\nabla}^{\text{blue}}(v, u)$. Since u lies in the white positive cone of v that does not contain w , by Lemma 18, we know that $d_{\nabla}(u, w) = d_{\nabla}(v, w) + \delta_{\nabla}^{\text{blue}}(v, u)$, and that $d_{\nabla}(v, u) \leq \delta_{\nabla}^{\text{blue}}(u, w)$. Therefore, we get $d_{\mathcal{S}}(u, w) \leq 9d_{\nabla}(u, w) + \delta_{\nabla}^{\text{blue}}(u, w)$ as desired. ◀

► **Lemma 20.** For any crossed blue edge $(u, w) \in \mathcal{D}$, $d_{\mathcal{S}}(u, w) \leq 3d_{\nabla}(u, w) + 9\delta_{\nabla}^{\min}(u, w)$.

Proof. Let (p, q) be a shortcut (p, q) crosses (u, w) . Since this shortcut is in the blue cone, it must have been added in \mathcal{S} replacing two white canonical edges incoming at u , namely (p, u) and (q, u) . As p and q are endpoints of the shortcut (p, q) , both blue edges (p, w) and (q, w) are uncrossed. Consequently, we have $d_{\mathcal{S}}(p, w) \leq 3d_{\nabla}(p, w)$ and $d_{\mathcal{S}}(q, w) \leq 3d_{\nabla}(q, w)$ by Lemma 14. Furthermore, by Lemmas 10 and 19 we know that both of the canonical edges (p, u) and (q, u) satisfy the inequalities $d_{\mathcal{S}}(p, u) \leq 9d_{\nabla}(p, u)$ and $d_{\mathcal{S}}(q, u) \leq 9d_{\nabla}(q, u)$. Since both of these canonical edges are incoming at u , we also know that one of them, say (p, u) , is not longer than $\delta_{\nabla}^{\min}(u, w)$, i.e., $d_{\nabla}(p, u) \leq \delta_{\nabla}^{\min}(u, w)$. Following the path via the point p , we bound the distance $d_{\mathcal{S}}(u, w)$ by $3d_{\nabla}(p, w) + 9d_{\nabla}(p, u)$, which in turn is bounded by

$3d_{\nabla}(u, w) + 9\delta_{\nabla}^{\min}(u, w)$ as stated in the lemma. \blacktriangleleft

► **Lemma 21.** *For any crossed white edge $(u, w) \in \mathcal{D}$, $d_{\mathcal{S}}(u, w) \leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{\min}(u, w)$.*

Proof. Since there are no shortcuts on the blue side of a white anchor, all white edges on the blue side are uncrossed. Therefore, we can assume that (u, w) is a white edge on the white side crossing a shortcut (p, q) in the same cone.

We consider two specific paths between u and w : *path A* which starts with the path for (p, w) as discussed in Lemma 19 and then follows the canonical path from p to u ; and *path B* which starts with the same path for (p, w) , and then takes the shortcut from p to q , and then follows the canonical path from q to u in the other direction. Considering both of the paths A and B, we bound $d_{\mathcal{S}}(p, w)$ by $9d_{\nabla}(p, w) + \delta_{\nabla}^{\text{blue}}(p, w)$ using Lemma 19. Also, for any blue canonical edge (x, y) on the canonical path between p and q , we know that $d_{\mathcal{S}}(x, y) \leq 3d_{\nabla}(x, y)$ by Lemmas 10 and 14. And, except for the first white canonical edge (p, r) , for any other white canonical edge (x', y') on this path, we know that $d_{\mathcal{S}}(x', y') \leq 2d_{\nabla}(x', y')$ by Lemmas 10 and 15, because either these white canonical edges themselves or their anchors are in \mathcal{S} . The first white canonical edge (p, r) and its anchor, however, may both be excluded from \mathcal{S} because of the shortcut (p, q) . Thus, we can only bound $d_{\mathcal{S}}(p, r)$ by the worse bound of $9d_{\nabla}(p, r) + \delta_{\nabla}^{\text{blue}}(p, r)$ using Lemmas 10 and 19. Using the fact that this cost has a much worse multiplicative constant than all other white canonical edges, it can be verified that the worst case happens when (p, r) is the only white canonical edge on the path from p to u .

We bound the total length of path A by $9d_{\nabla}(p, w) + \delta_{\nabla}^{\text{blue}}(p, w)$ for the edge (p, w) , by $9d_{\nabla}(p, r) + \delta_{\nabla}^{\text{blue}}(p, r) \leq 10d_{\nabla}(p, r)$ for the white canonical edges between p and u , and by $3d_{\nabla}(r, u)$ for the blue canonical edges between p and u . Also, by using arguments analogous to those used in the proof of Lemma 18, we observe that $d_{\nabla}(p, r) \leq d_{\nabla}(p, u)$ and $d_{\nabla}(r, u) \leq \delta_{\nabla}^{\text{white}}(p, u)$. Therefore, the total cost of path A is bounded by

$$d_{\mathcal{S}}(u, w) \leq 9d_{\nabla}(p, w) + \delta_{\nabla}^{\text{blue}}(p, w) + 13d_{\nabla}(p, u) - 3\delta_{\nabla}^{\text{blue}}(p, u).$$

Since (p, u) is a white canonical edge in a white cone, by Lemma 18, we can substitute $d_{\nabla}(p, u) = \delta_{\nabla}^{\text{blue}}(u, w) - \delta_{\nabla}^{\text{blue}}(p, w)$ and $\delta_{\nabla}^{\text{blue}}(p, u) = d_{\nabla}(u, w) - d_{\nabla}(p, w)$. Hence, we get

$$\begin{aligned} d_{\mathcal{S}}(u, w) &\leq 12d_{\nabla}(p, w) - 12\delta_{\nabla}^{\text{blue}}(p, w) + 13\delta_{\nabla}^{\text{blue}}(u, w) - 3d_{\nabla}(u, w) \\ &= 12\delta_{\nabla}^{\text{white}}(p, w) + 10\delta_{\nabla}^{\text{blue}}(u, w) - 3\delta_{\nabla}^{\text{white}}(u, w) \\ &= 12\delta_{\nabla}^{\text{white}}(p, w) + 10d_{\nabla}(u, w) - 13\delta_{\nabla}^{\text{white}}(u, w) \\ &= 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{\text{white}}(u, w) + 12\delta_{\nabla}^{\text{white}}(p, w) - 23\delta_{\nabla}^{\text{white}}(u, w). \end{aligned} \quad (1)$$

By using the fact that $\delta_{\nabla}^{\text{white}}(p, w) \leq d_{\nabla}(u, w)$ and rearranging (1) we get:

$$\begin{aligned} d_{\mathcal{S}}(u, w) &\leq 10\delta_{\nabla}^{\text{blue}}(u, w) + 20\delta_{\nabla}^{\text{white}}(u, w) + 10d_{\nabla}(u, w) + 2\delta_{\nabla}^{\text{white}}(p, w) - 23\delta_{\nabla}^{\text{white}}(u, w) \\ &\leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{\text{blue}}(u, w) + 2\delta_{\nabla}^{\text{white}}(p, w) - 3\delta_{\nabla}^{\text{white}}(u, w). \end{aligned}$$

Hence, we conclude that $d_{\mathcal{S}}(u, w) \leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{\min}(u, w)$ if $2\delta_{\nabla}^{\text{white}}(p, w) \leq 3\delta_{\nabla}^{\text{white}}(u, w)$. Therefore, we assume $3\delta_{\nabla}^{\text{white}}(u, w) < 2\delta_{\nabla}^{\text{white}}(p, w) \leq 2d_{\nabla}(u, w) = 2\delta_{\nabla}^{\text{white}}(u, w) + 2\delta_{\nabla}^{\text{blue}}(u, w)$, and derive $\delta_{\nabla}^{\text{white}}(u, w) < 2\delta_{\nabla}^{\text{blue}}(u, w)$.

Next, we bound the total length of path B by $9d_{\nabla}(p, w) + \delta_{\nabla}^{\text{blue}}(p, w)$ for the edge (p, w) , by $d_{\nabla}(p, q)$ for the shortcut, by $2d_{\nabla}(u, q)$ for the white canonical edges between p and u , and by $3d_{\nabla}(u, q)$ for the blue canonical edges between p and u ; thus in total by

$$d_{\mathcal{S}}(u, w) \leq 9d_{\nabla}(p, w) + \delta_{\nabla}^{\text{blue}}(p, w) + d_{\nabla}(p, q) + 5d_{\nabla}(u, q).$$

By the triangle inequality, we have $d_{\nabla}(p, q) \leq d_{\nabla}(p, u) + d_{\nabla}(u, q)$ and using once again that $d_{\nabla}(p, u) = \delta_{\nabla}^{blue}(u, w) - \delta_{\nabla}^{blue}(p, w)$, we bound the total cost of path B by

$$d_S(u, w) \leq 9d_{\nabla}(p, w) + \delta_{\nabla}^{blue}(u, w) + 6d_{\nabla}(u, q). \quad (2)$$

We further bound the cost of path B if $d_{\nabla}(u, w) \geq d_{\nabla}(q, w)$, in which case we observe $d_{\nabla}(u, q) \leq \delta_{\nabla}^{white}(u, w)$. Using the inequality $\delta_{\nabla}^{white}(u, w) < 2\delta_{\nabla}^{blue}(u, w)$, we rearrange (2)

$$\begin{aligned} d_S(u, w) &\leq 9d_{\nabla}(p, w) + \delta_{\nabla}^{blue}(u, w) + 6d_{\nabla}(u, q) \\ &\leq 9d_{\nabla}(p, w) + \delta_{\nabla}^{blue}(u, w) + 6\delta_{\nabla}^{white}(u, w) \\ &\leq 10d_{\nabla}(u, w) + 5\delta_{\nabla}^{white}(u, w) \\ &\leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{blue}(u, w) \\ &\leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{\min}(u, w). \end{aligned}$$

Therefore, we consider the remaining case $d_{\nabla}(u, w) < d_{\nabla}(q, w)$. In this case, by Lemma 18, we get $\delta_{\nabla}^{blue}(u, q) = d_{\nabla}(q, w) - d_{\nabla}(u, w)$, and we also observe that $\delta_{\nabla}^{white}(u, q) \leq \delta_{\nabla}^{white}(u, w)$. Combining them, we get $d_{\nabla}(u, q) \leq d_{\nabla}(q, w) - \delta_{\nabla}^{blue}(u, w)$, and rearrange (2) we get

$$d_S(u, w) \leq 9d_{\nabla}(p, w) + 6d_{\nabla}(q, w) - 5\delta_{\nabla}^{blue}(u, w). \quad (3)$$

Furthermore, using the earlier assumption that $3\delta_{\nabla}^{white}(u, w) < 2\delta_{\nabla}^{white}(p, w)$, and the fact that $\delta_{\nabla}^{white}(q, w) \geq 0$, we deduce that $3\delta_{\nabla}^{blue}(u, q) < 2\delta_{\nabla}^{blue}(p, u) + \delta_{\nabla}^{blue}(u, q)$. Using Lemma 18 on all the three terms, we obtain $d_{\nabla}(q, w) < 3d_{\nabla}(u, w) - 2d_{\nabla}(p, w)$. We then rearrange (3)

$$\begin{aligned} d_S(u, w) &\leq 9d_{\nabla}(p, w) + 18d_{\nabla}(u, w) - 12d_{\nabla}(p, w) - 5\delta_{\nabla}^{blue}(u, w) \\ &= 10d_{\nabla}(u, w) + 8\delta_{\nabla}^{white}(u, w) + 3\delta_{\nabla}^{blue}(u, w) - 3d_{\nabla}(p, w) \\ &\leq 10d_{\nabla}(u, w) + 8\delta_{\nabla}^{white}(u, w) + 3\delta_{\nabla}^{blue}(u, w) - 3\delta_{\nabla}^{white}(p, w) \\ &\leq 10d_{\nabla}(u, w) + \frac{7}{2}\delta_{\nabla}^{white}(u, w) + 3\delta_{\nabla}^{blue}(u, w) \\ &\leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{blue}(u, w). \end{aligned}$$

Therefore, what remains to be proven is $d_S(u, w) \leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{\min}(u, w)$ for the case $\delta_{\nabla}^{\min}(u, w) = \delta_{\nabla}^{white}(u, w)$. To prove this case, using (1), we will assume the stronger inequality $7\delta_{\nabla}^{white}(u, w) < 4\delta_{\nabla}^{white}(p, w)$, and using arguments similar to those used in the above paragraph, we get $3d_{\nabla}(q, w) < 7d_{\nabla}(u, w) - 4d_{\nabla}(p, w)$. Once again, we rearrange (3)

$$\begin{aligned} d_S(u, w) &\leq 9d_{\nabla}(p, w) + 14d_{\nabla}(u, w) - 8d_{\nabla}(p, w) - 5\delta_{\nabla}^{blue}(u, w) \\ &= 9d_{\nabla}(u, w) + d_{\nabla}(p, w) + 5\delta_{\nabla}^{white}(u, w) \\ &\leq 10d_{\nabla}(u, w) + 5\delta_{\nabla}^{white}(u, w). \end{aligned}$$

Proving $d_S(u, w) \leq 10d_{\nabla}(u, w) + 10\delta_{\nabla}^{\min}(u, w)$ for all of the cases, we conclude that the statement of the lemma is correct. ◀

By Lemmas 14, 19, 20, and 21 we have that \mathcal{S} is a 20-spanner of \mathcal{D} . Since \mathcal{D} is a 2-spanner of \mathcal{C} ([11]) it follows that \mathcal{S} is a spanner of \mathcal{C} with stretch factor at most 40. We prove, however, a much better bound of 20 next.

► **Lemma 22.** *For any two points $p, q \in \mathcal{P}$, $d_{\mathcal{S}}(p, q)$ is bounded by $20d_{\nabla}(p, q)$.*

Proof. We prove the lemma by first constructing, in \mathcal{D} , a monotone path π between p and q that lies inside $\nabla(p, q)$. We then consider the path π' in \mathcal{S} obtained by replacing every edge of π not in \mathcal{S} with a short path in \mathcal{S} .

We define the path π between p and q consisting of k edges in \mathcal{D} using a sequence of pairs of points $\{p, q\} = \{p_0, q_0\}, \{p_1, q_1\}, \dots, \{p_k, q_k\}$ such that any two consecutive pairs of points $\{p_{i-1}, q_{i-1}\}$ and $\{p_i, q_i\}$ satisfy exactly one of the equations $p_i = p_{i-1}$ and $q_i = q_{i-1}$ and the equation that is not satisfied describes the i^{th} edge. If $p_i \neq p_{i-1}$, then the i^{th} edge is $(p_{i-1}, p_i) \in \mathcal{D}$, otherwise the i^{th} edge is $(q_{i-1}, q_i) \in \mathcal{D}$. We define this sequence recursively for the next pair of points $\{p_{i+1}, q_{i+1}\}$ by first identifying which of the points p_i and q_i lie in the other's positive cone. If q_i lies in the positive cone of p_i , then we define $q_{i+1} = q_i$ and $p_{i+1} = r$ such that $(p_i, r) \in \mathcal{D}$, noting that by definition of \mathcal{D} , r is the unique such point in the positive cone of p_i that contains q_i . Otherwise, if p_i lies in a positive cone of q_i , then we define $p_{i+1} = p_i$ and $q_{i+1} = r'$ such that $(q_i, r') \in \mathcal{D}$. We stop when $p_k = q_k$.

We prove inductively that the aforementioned path π lies within $\nabla(p, q)$. For the base case, clearly the path consisting of the only edge in the sequence $\{p_{k-1}, q_{k-1}\}, \{p_k, q_k\}$ lies within $\nabla(p_{k-1}, q_{k-1})$. For the inductive step, assuming that the path for the sequence $\{p_i, q_i\}, \dots, \{p_k, q_k\}$ lies within $\nabla(p_i, q_i)$, we show that the path for the sequence $\{p_{i-1}, q_{i-1}\}, \{p_i, q_i\}, \dots, \{p_k, q_k\}$ lies within $\nabla(p_{i-1}, q_{i-1})$. First, we observe that in either case that the first edge is (p_{i-1}, p_i) or (q_{i-1}, q_i) , it lies within $\nabla(p_{i-1}, q_{i-1})$ by definition. Finally, we observe that $\nabla(p_i, q_i)$, hence the rest of the path lies within $\nabla(p_{i-1}, q_{i-1})$ as well. Therefore, we prove the inductive step.

We then prove that all of the edges of the form (p_i, p_{i+1}) lie in the same corresponding positive cones of their respective points p_i . More specifically, we prove by induction on the sequence of such edges $e_0 = (p = p_{i_0-1}, p_{i_0}), e_1 = (p_{i_0} = p_{i_1-1}, p_{i_1}), \dots, e_\ell = (p_{i_\ell-1} = p_{i_\ell-1}, p_{i_\ell} = p_k)$, that $p_{i_0}, p_{i_1}, \dots, p_{i_\ell}$ lie in the same corresponding cones of $p_{i_0-1}, p_{i_1-1}, \dots, p_{i_\ell-1}$ respectively.

The base case follows trivially, and for the inductive step we assume for edges $e_0, e_1, \dots, e_\lambda$ that $p_{i_0}, p_{i_1}, \dots, p_{i_\lambda}$ lie in the same corresponding cones of $p_{i_0-1}, p_{i_1-1}, \dots, p_{i_\lambda-1}$ respectively. For the inductive step we need to prove for the edge $e_{\lambda+1}$ that $p_{i_{\lambda+1}}$ lies in the same corresponding cone of $p_{i_{\lambda+1}-1} = p_{i_\lambda}$. We have already proven that the edge $e_{\lambda+1}$ lies within $\nabla(p_{i_\lambda}, q_{i_\lambda})$. Also, by definition of \mathcal{D} , we know that $\nabla(p_{i_\lambda-1}, p_{i_\lambda})$ is empty of points of \mathcal{P} in its interior. Then we conclude the inductive proof by observing that the only positive cone that can possibly include $e_{\lambda+1}$ at p_{i_λ} is part of the same corresponding cone of p_{i_λ} . Having proven this critical property about this path, we denote it by π_p and refer to it as one of the two *branches*, where the other *branch* π_q is defined analogously using points q_0, q_1, \dots, q_k . We conclude that the path π consisting of these two branches π_p and π_q is monotone.

Finally, we prove the claimed bound on the length of the path π' between p and q . Because the constants in Lemma 21 are the largest among Lemmas 14, 19, 20, and 21, the worst case happens when q lies in the white positive cone of p and π is white monotone. Let z be the blue vertex of $\nabla(p, q)$. By Lemma 6: the projections of all edges of π onto zp (resp. zq) do not overlap, are contained within $[zp]$ (resp. zq), $|zp| = d_{\nabla}(p, q)$, and $|zq| = \delta_{\nabla}^{\text{blue}}(p, q)$. In the worst case, each edge of π is crossed, and Lemma 21 applies to reconstruct each edge of π . Therefore, the length of π' can be upper bounded by $10d_{\nabla}(p, q) + 10\delta_{\nabla}^{\min}(p, q) \leq 10d_{\nabla}(p, q) + 10\delta_{\nabla}^{\text{blue}}(p, q) \leq 20d_{\nabla}(p, q)$. ◀

► **Theorem 23.** *\mathcal{S} is a plane spanner of \mathcal{C} with maximum degree 4 and stretch factor at most 20 that can be constructed in $\mathcal{O}(n \log n)$ time.*

Proof. The properties of \mathcal{S} regarding planarity, maximum degree bound, and stretch factor are proven by Lemmas 11, 13, and 22 respectively. The TD-Delaunay triangulation \mathcal{D} can be constructed in $\mathcal{O}(n \log n)$ time [11]. Given \mathcal{D} and the fact that it is planar, \mathcal{S} can be constructed in $\mathcal{O}(n \log n)$ time: sorting the anchors takes $\mathcal{O}(n \log n)$ time, and adding edges to \mathcal{S} can be done in $\mathcal{O}(n)$ time. ◀

6 Tight degree bound for points in convex position

In this section, we show that if the set \mathcal{P} of points is in convex position, then the *same* spanner \mathcal{S} constructed in the previous section has maximum degree at most 3. Therefore, for any set of points \mathcal{P} in convex position, there exists a plane spanner of \mathcal{C} of maximum degree at most 3. We also show in this section that 3 is a lower bound on the maximum degree of plane spanners of \mathcal{C} for point-sets in convex position. The preceding implies that 3 is a tight bound on the maximum degree of plane spanners of \mathcal{C} for point-sets in convex position, and this completely and satisfactorily answers the question about the maximum degree of plane geometric spanners of \mathcal{C} , in the case when \mathcal{P} is in convex position.

► **Proposition 24.** Let \mathcal{P} be a set of points in convex position in the plane, and let \mathcal{S} be the spanner of \mathcal{C} constructed as described in Section 4. Then the maximum degree of \mathcal{S} is at most 3.

Proof. Let $p \in \mathcal{P}$. It suffices to show that the degree of p in \mathcal{S} is at most 3. To this effect, we show that at most 3 edges incident to p could be added in steps 1 – 6 of the construction of \mathcal{S} . Since \mathcal{P} is in convex position, there exists a support line, D_p , passing through p such that all the points of \mathcal{P} lie in one closed half plane, H , of the two half planes delimited by D_p [17]. Observe that – by the construction of \mathcal{S} – each of the two blue sectors of p contains at most one edge incident to p in \mathcal{S} . In Lemma 12, we showed that each of the at most 4 edges incident to a point $p \in \mathcal{S}$ is charged by the charging scheme (see Section 5) to one of the 4 sectors LB_p , RB_p , UW_p , and LW_p , such that each of the 4 sectors is charged with at most one edge. In this charging scheme, a blue anchor at p is charged to the blue sector containing it, and a white edge incident to p is either charged to a blue sector at p (and in such case the blue sector does not contain a blue anchor), or to the white sector containing the edge. We distinguish the following cases, based on the angle, α , that D_p makes with the positive x -axis:

- Case 1. $0 \leq \alpha \leq \pi/3$ (D_p passes through the two blue sectors at p). In this case the half plane H entirely contains one of the two white sectors at p . Suppose that H contains UW_p ; the case is symmetric if H contains LW_p . Since (1) each of the three sectors LB_p , UW_p , and RB_p is charged with at most one edge incident to p , (2) each of LB_p and RB_p contains at most one blue anchor incident to p , and (3) each edge incident to p in UW_p is either charged to UW_p , LB_p , or RB_p , it follows that the number of edges incident to p in H , and hence in \mathcal{S} , is at most 3.
- Case 2. $\pi/3 \leq \alpha \leq 2\pi/3$ (D_p passes through the two red cones at p). Assume, to get a contradiction, that there are 4 edges incident to p in H . Suppose first that H entirely contains RB_p , and hence is disjoint from LB_p . In this case the part of UW_p in H is contained in a positive cone at p , and hence, there is at most one edge incident to p in \mathcal{S} that lies in $H \cap UW_p$. Since each of the 4 sectors at p is charged with at most one

edge, it follows from the preceding that there is a white edge e incident to p in $H \cap LW_p$ that is charged to LB_p . This could only happen if e , a white edge, is a canonical edge in a blue cone, added according to step 3 in the spanner construction, and is charged to LB_p according to step 3 of the charging scheme. This, however, implies that LB_p is not empty (a white canonical edge $e \in LW_p$ in a blue cone could exist only if RB_p contains points of \mathcal{P}), contradicting that all points of \mathcal{P} lie in H . Suppose now that H entirely contains LB_p , and hence is disjoint from RB_p . By the same token as above, there must exist a white canonical edge, e , incident to p and lying in a white sector at p , that is charged to LB_p . By the charging scheme, the edge e cannot lie in UW_p because in such case e would not be charged to UW_p . Therefore, e must lie in LW_p . But then e must lie in the negative cone adjacent to p in LW_p , which is (the cone) disjoint from H , again contradicting that H contains all points of \mathcal{P} .

- Case 3. $2\pi/3 \leq \alpha \leq \pi$ (D_p passes through the two green cones at p). The proof is analogous to that of Case 2 above. ◀

► **Proposition 25.** For any constant $\rho \geq 1$, there exists a point-set \mathcal{P} in convex position such that any plane spanner of \mathcal{C} of maximum degree at most 2 has stretch factor $> \rho$.

Proof. Let $\rho \geq 1$ be a given constant. Choose an integer $b > \rho$, and an integer $N > 3(\rho \cdot b + 1)$. Consider an orthogonal rectangle of vertical dimension $a = N - 1$ and horizontal dimension b . Let $n = 2N$, and let $\mathcal{P} = \{p_1, \dots, p_N, q_1, \dots, q_N\}$ be a set of n points placed on the rectangle as follows. Points p_1, \dots, p_N are placed on one vertical side of the rectangle such that $|p_i p_{i+1}| = 1$, for $i = 1, \dots, N - 1$, so that p_1 and p_N end up on the two vertices of the vertical side of the rectangle. Points q_1, \dots, q_N are placed symmetrically on the other vertical side of the rectangle so that the $p_i q_i$'s, $i = 1, \dots, N$, are all parallel. We note that one can choose \mathcal{P} to be in convex position while respecting the standard general position assumptions (no 3 points on a line, etc.) by slightly modifying the set of points chosen above (e.g., rotating each of $p_i p_{i+1}$ and $q_i q_{i+1}$, $i = 1, \dots, N$, slightly but increasingly inwards (*i.e.*, towards the interior of the rectangle), and slightly modifying the choice of the parameters a, b, N). However, we decided to go with the above configuration for clarity and ease of presentation.

Suppose, to get a contradiction, that there is a plane spanner S of \mathcal{C} of maximum degree at most 2 and stretch factor ρ . Since S is connected, S is either a Hamiltonian path or a Hamiltonian cycle. Without loss of generality, we will assume in what follows that S is a Hamiltonian cycle, as the proof in the other case is simpler. Let $L = \{p_1, \dots, p_N\}$ and $R = \{q_1, \dots, q_N\}$. We distinguish the following two cases:

- Case 1. The subgraph of S induced by either L or R is disconnected. Suppose that the subgraph of S induced by L , S_L , is disconnected; the proof follows by symmetry in the other case. Then there must exist $i \in \{1, \dots, N - 1\}$ such that the two consecutive points p_i, p_{i+1} belong to two different connected components of S_L . Therefore, any path between p_i and p_{i+1} in S must contain an edge $p_r q_s$, $r, s \in \{1, \dots, N\}$. Since the distance between any point from L and any point from R is at least b , we have $|p_r q_s| \geq b$. It follows that the length of any path in S between p_i and p_{i+1} is at least $b > \rho = \rho \cdot |p_i p_{i+1}|$, which contradicts that S is a spanner of stretch factor ρ .
- Case 2. Each of the two subgraphs S_L and S_R of S , induced by L and R , respectively, is connected. Then each of S_L and S_R must be a Hamiltonian path on L and R , respectively. Let p_r and p_s , $r, s \in \{1, \dots, N\}$, be the points of degree 1 in S_L , and $q_{r'}, q_{s'}$, $r', s' \in \{1, \dots, N\}$, be those of degree 1 in S_R . Since S is a Hamiltonian cycle, S consists of the edges in the

Hamiltonian path S_L , the edges in the Hamiltonian path S_R , plus a matching between $\{p_r, p_s\}$ and q'_r, q'_s , say $p_r p_{r'}$ and $q_s q_{s'}$. Since there are N points in L , there must exist a point $p_i \in L$, $i \in \{1, \dots, N\}$, such that the number of edges on each of the two subpaths between p_i and p_r and between p_i and p_s in S_L , is at least $N/3 - 1 > \rho \cdot b$ by the choice of N . Consider point $q_i \in R$. Any path between p_i and q_i in S must contain either the subpath of S_L between p_i and p_r , or the subpath of S_L between p_i and p_s , and hence must contain more than $\rho \cdot b$ edges of S_L . Since each edge of S_L has length at least 1, the length of any path between p_i and q_i in S is more than $\rho \cdot b = \rho \cdot |p_i q_i|$. This again contradicts the assumption that S has stretch factor ρ , and completes the proof. ◀

We conclude with the following theorem:

► **Theorem 26.** *The constant 3 is a tight bound on the maximum degree of geometric plane spanners of \mathcal{C} for point-sets in convex position.*

References

- 1 N. Bonichon, C. Gavoille, N. Hanusse, and D. Ilcinkas. Connections between theta-graphs, delaunay triangulations, and orthogonal surfaces. In *Proceedings of the 36th International Workshop on Graph Theoretic Concepts in Computer Science*, volume 6410 of *Lecture Notes in Computer Science*, pages 266–278, 2010.
- 2 N. Bonichon, C. Gavoille, N. Hanusse, and L. Perković. Plane spanners of maximum degree six. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 6198 of *Lecture Notes in Computer Science*, pages 19–30. Springer, 2010.
- 3 N. Bonichon, C. Gavoille, N. Hanusse, and L. Perković. The stretch factor of L_1 - and L_∞ -Delaunay triangulations. In *Proceedings of the 20th Annual European Symposium on Algorithms (ESA)*, volume 7501 of *Lecture Notes in Computer Science*, pages 205–216. Springer, 2012.
- 4 N. Bonichon, I. Kanj, L. Perkovic, and G. Xia. There are plane spanners of degree 4 and moderate stretch factor. *Discrete & Computational Geometry*, 53(3):514–546, 2015.
- 5 P. Bose, P. Carmi, and L. Chaitman-Yerushalmi. On bounded degree plane strong geometric spanners. *J. Discrete Algorithms*, 15:16–31, 2012.
- 6 P. Bose, P. Carmi, S. Collette, and M. Smid. On the stretch factor of convex delaunay graphs. *Journal of Computational Geometry*, 1(1):41–56, 2010.
- 7 P. Bose, J. Gudmundsson, and M. Smid. Constructing plane spanners of bounded degree and low weight. *Algorithmica*, 42(3-4):249–264, 2005.
- 8 P. Bose, P. Morin, I. Stojmenović, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *WIRELESS NETWORKS*, 7(6):609–616, 2001.
- 9 P. Bose, M. Smid, and D. Xu. Delaunay and diamond triangulations contain spanners of bounded degree. *International Journal of Computational Geometry and Applications*, 19(2):119–140, 2009.
- 10 L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proceedings of the Second Annual Symposium on Computational Geometry (SoCG)*, pages 169–177, 1986.
- 11 L. P. Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer and System Sciences*, 39(2):205–219, 1989.
- 12 G. Das and P.J. Heffernan. Constructing degree-3 spanners with other sparseness properties. *Int. J. Found. Comput. Sci.*, 7(2):121–136, 1996.

- 13 D. Dobkin, S. Friedman, and K. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete & Computational Geometry*, 5(4):399–407, December 1990. doi:10.1007/BF02187801.
- 14 I. Kanj and L. Perković. On geometric spanners of Euclidean and unit disk graphs. In *In Proceedings of the 25th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume hal-00231084, pages 409–420. HAL, 2008.
- 15 J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7(1):13–28, 1992.
- 16 X.-Y. Li and Y. Wang. Efficient construction of low weight bounded degree planar spanner. *International Journal of Computational Geometry and Applications*, 14(1-2):69–84, 2004.
- 17 F. Preparata and M. Shamos. *Computational Geometry - An Introduction*. Springer, 1985.
- 18 J. Salowe. Euclidean spanner graphs with degree four. *Discrete Applied Mathematics*, 54(1):55–66, 1994.
- 19 Y. Wang and Xiang-Yang L. Localized construction of bounded degree and planar spanner for wireless ad hoc networks. *Mobile Networks and Applications*, 11(2):161–175, 2006.
- 20 G. Xia. The stretch factor of the Delaunay triangulation is less than 1.998. *SIAM J. Comput.*, 42(4):1620–1659, 2013.